



ADVANCED COMPUTING RESEARCH LABORATORY

CHEP 2012 May 20 – 25, 2012 Lennart Johnsson



Computing Technology Future

Lennart Johnsson, University of Houston, Houston, TX







ADVANCED COMPUTING RESEARCH LABORATOR

Outline

CHEP 2012 May 20 – 25, 2012 Lennart Johnsson



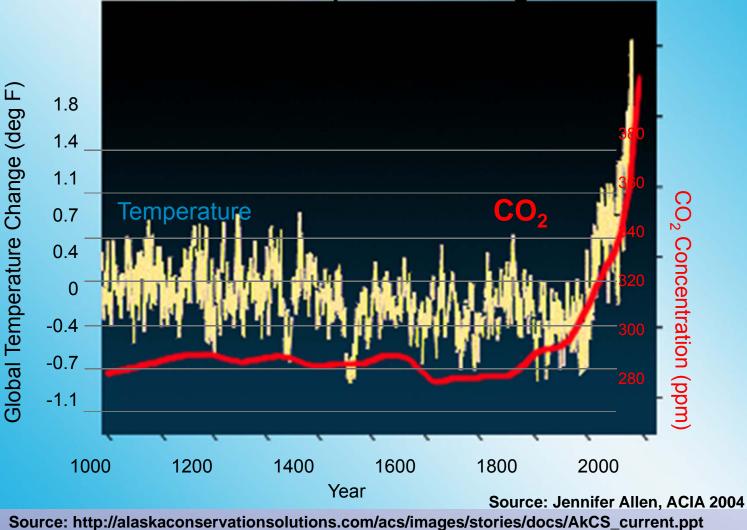
- Energy of computation and its impact
- What kind of architecture?
- Some possible approaches





ADVANCED COMPUTING RESEARCH LABORATOR

1000 Years of CO₂ and Global Temperature Change

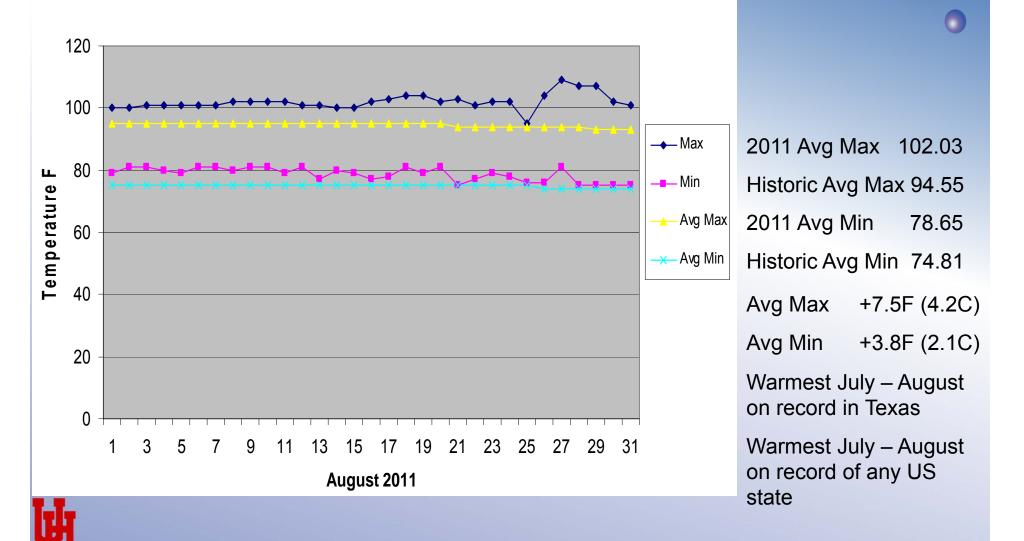


CHEP 2012 May 20 – 25, 2012 Lennart Johnsson

0



Houston August 2011 Daily Temperatures









ADVANCED COMPUTING RESEARCH LABORATOR

Severe Weather



Hurricanes:

For 1925 - 1995 the US cost was \$5 billion/yr for a total of 244 landfalls. But, hurricane Andrew alone caused damage in excess of \$27 billion.

The US loss of life has gone down to <20/yr typically. The Galveston "Great Hurricane" year 1900 caused over 6,000 deaths.

Since 1990 the number of landfalls each year is increasing.

Warnings and emergency response costs on average \$800 million/yr. Satellites, forecasting efforts and research cost \$200 – 225 million/yr.







Tornados



http://en.wikipedia.org/wiki/ File:Dimmit_Sequence.jpg



http://www.miapearlman.com/ images/tornado.jpg



www.drjudywood.com/.../spics/ tornado-760291.jpg

http://g.imagehost.org/0819 /tornado-damage.jpg

http://www.crh.noaa.gov/mkx/ document/tor/images/tor060884/ damage-1.jpg

















Russia may have lost 15,000 lives already, and \$15 billion, or 1% of GDP, according to Bloomberg.

The smog in Moscow is a driving force behind the fires' deadly impact, with 7000 being killed already in the city. Aug 10, 2010

Iddfires Wildfires

In a single week, San Diego County wildfires killed 16 people, destroyed nearly 2,500 homes and burned nearly 400,000 acres. Oct 2003

http://legacy.signonsandiego.com/ news/fires/weekoffire/images/mainimage4.jpg



Russia Wildfires 2010

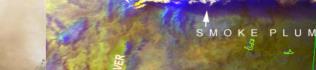
Russian-fires-control.jpg

http://topnews.in/law/files/

http://msnbcmedia1.msn.com/j/ MSNBC/Components/Photo/_new/ 100810-russianFire-vmed-218p. grid-6x2.jpg



Los Alamos Forest Fires NOAA-15 AVHRR HRPT Multi-channel False Color Image May 11, 2000 @ 0122 UTC



Fires



http://www.tolerance.ca/image/ photo 1281943312664-2-0 94181 G.jpc

img.ibtimes.com/www/data/images/full/2010/08/



April 30 – May 7, 2010 TN, KY, MS 31 deaths. Nashville Mayor Karl Dean estimates the damage from weekend flooding could easily top \$1 billion.





Floods

UK June – July 2007 13 deaths more than 1 million affected cost about £6 billion





CHEP 2012 May 20 – 25, 2012 Lennart Johnsson



China,Bloomberg Aug 17,2010 1450 deaths through Aug 6 Aug 7 1254 killed in mudslide with 490 missing

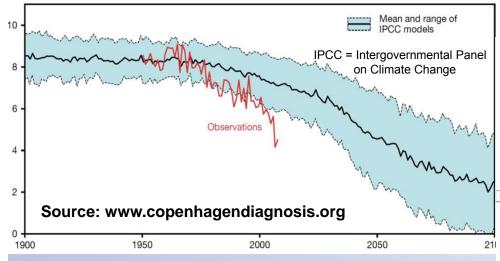






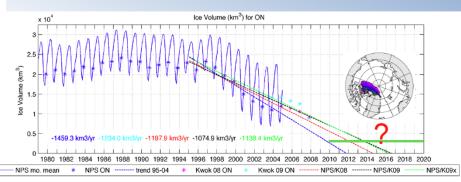
ADVANCED COMPUTING RESEARCH LABORATOR

Arctic Summer Ice Melting Accelerating



Ice area millions km², September minimum

Ice Volume km³



Observational estimates (cyan / purple stars):

- Obs Fall (ON) '07 volume <9000 km³, ~20% uncertainty,

- Negative volume trends: 1197 - 1234 km³/yr

- Combined (95-07) model / data linear volume trend projects ice-free fall by 2016

- Same trend with extended K09 (assuming the constant ON volume for 07 - 09)

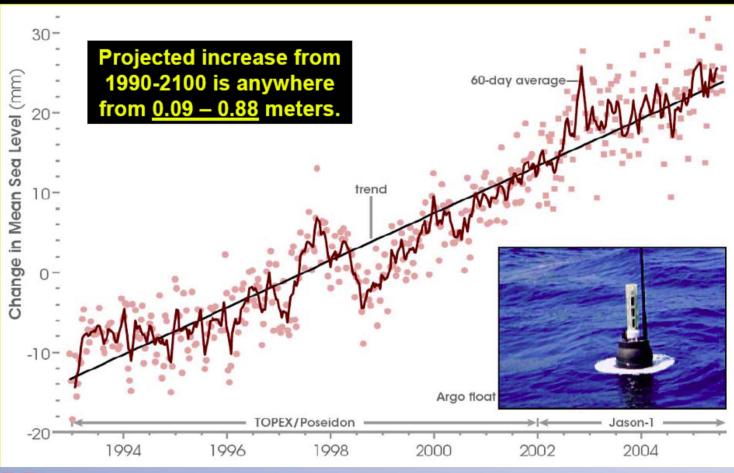
- Some (?) sea ice will remain beyond due to increased ridging of thinner ice

- Uncertainty (95-07) is ±3yrs and not all volume must disappear Kwok et al., JGR 2009, Kwok & Cunnigham, JGR 2008





Sea level has increased about 3 mm/yr between 1993 and 2005



1/3rd due to melting glaciers

2/3rd due expansion from warming oceans

Source: Trenberth, NCAR 2005



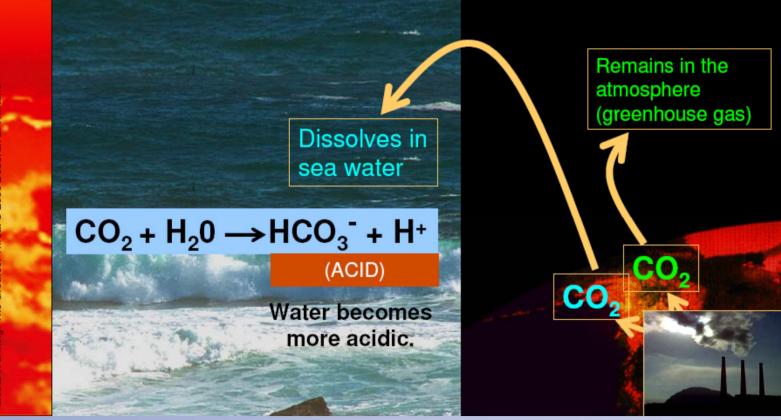
Source: lskhaq lskandar, http://www.jsps.go.jp/j-sdialogue/2007c/data/52_dr_iskander_02.pdf



Global Cataclysmic Concerns

Ocean Acidification

Over the last 200 years, about 50% of all CO₂ produced on earth has been absorbed by the ocean. (Royal Society 6/05)



Source: http://alaskaconservationsolutions.com/acs/images/stories/docs/AkCS_current.ppt



Global Cataclysmic Concerns

Ocean Acidification

Animals with calcium carbonate shells -- corals, sea urchins, snails, mussels, clams, certain plankton, and others -- have trouble building skeletons and shells can even begin to dissolve. "Within decades these shell-dissolving conditions are projected to be reached and to persist throughout most of the year in the polar oceans." (Monaco Declaration 2008)



- Pteropods (an important food source for salmon, cod, herring, and pollock) likely not able to survive at CO₂ levels predicted for 2100 (600ppm, pH 7.9) (Nature 9/05)
- Coral reefs at serious risk; doubling CO₂, stop growing and begin dissolving (GRL 2009)
- Larger animals like squid may have trouble extracting oxygen
- Food chain disruptions



Source: http://alaskaconservationsolutions.com/acs/images/stories/docs/AkCS_current.ppt



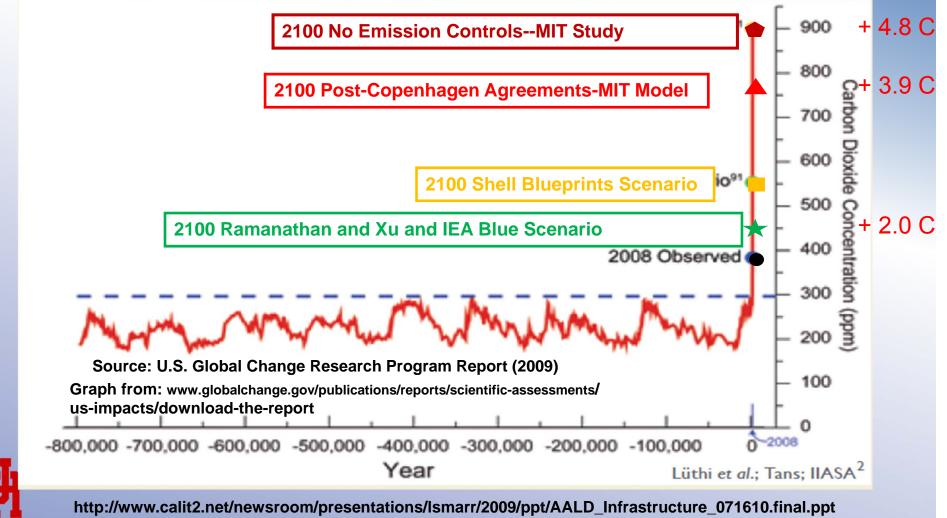




NCED COMPLITING RESEA

CHEP 2012 May 20 – 25, 2012 Lennart Johnsson

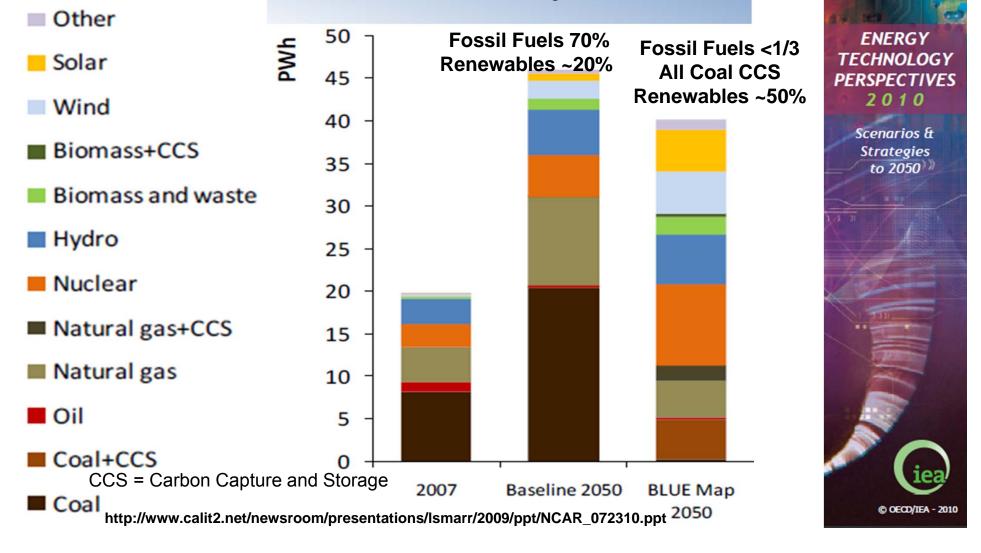
Atmospheric CO₂ Levels for Last 800,000 Years and Several Projections for the 21st Century







IEA Blue Map Requires Massive Decarbonising of the Electricity Sector







ADVANCED COMPUTING RESEARCH LABORATORY

CHEP 2012 May 20 – 25, 2012 Lennart Johnsson

ICT impact on CO₂ emissions*

- It is estimated that the ICT industry alone produces CO₂ emissions that is equivalent to the carbon output of the entire aviation industry. Direct emissions of Internet and ICT amounts to 2-3% of world emissions
- ICT emissions growth fastest of any sector in society; expected to double every 4 to 6 years with current approaches
- One small computer server generates as much carbon dioxide as a SUV with a fuel efficiency of 15 miles per gallon

*An Inefficient Tuth: http://www.globalactionplan.org.uk/event_detail.aspx?eid=2696e0e0-28fe-4121-bd36-3670c02eda49



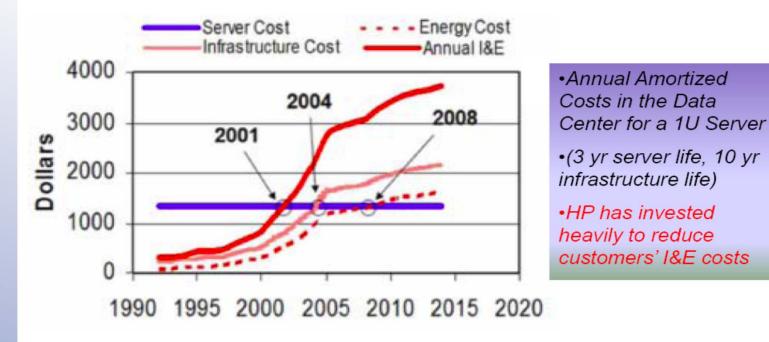




ADVANCED COMPUTING RESEARCH LABORATOR

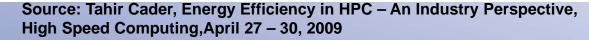
Evolution of Data Center Energy Costs (US)

The Cost to Power & Cool a Server Has Exceeded the Cost of the Server...



Source: Belady, C., 2007, "In the Data Center, Power and Cooling Costs More than IT Equipment it Supports", Electronics Cooling Magazine (Feb issue).

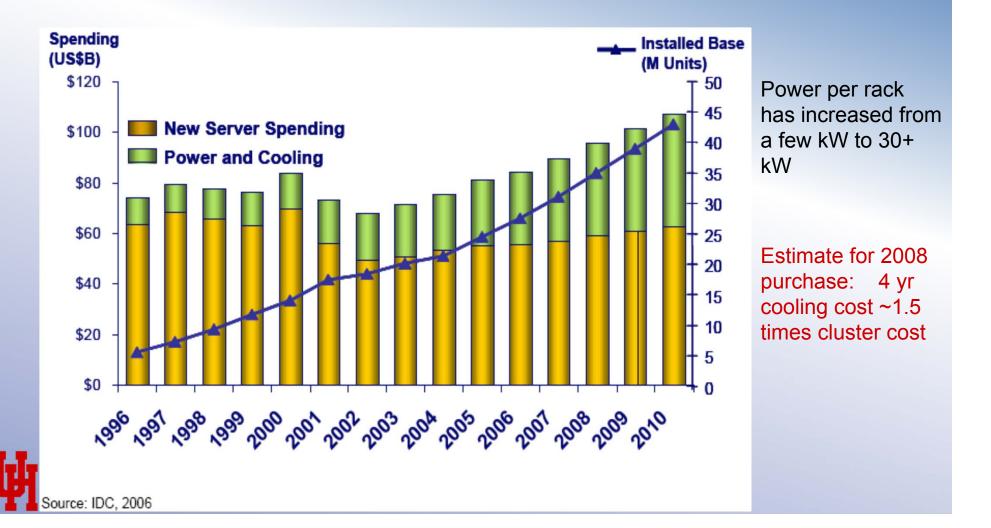








Worldwide Server Installed Base, New Server Spending, and Power and Cooling Expense

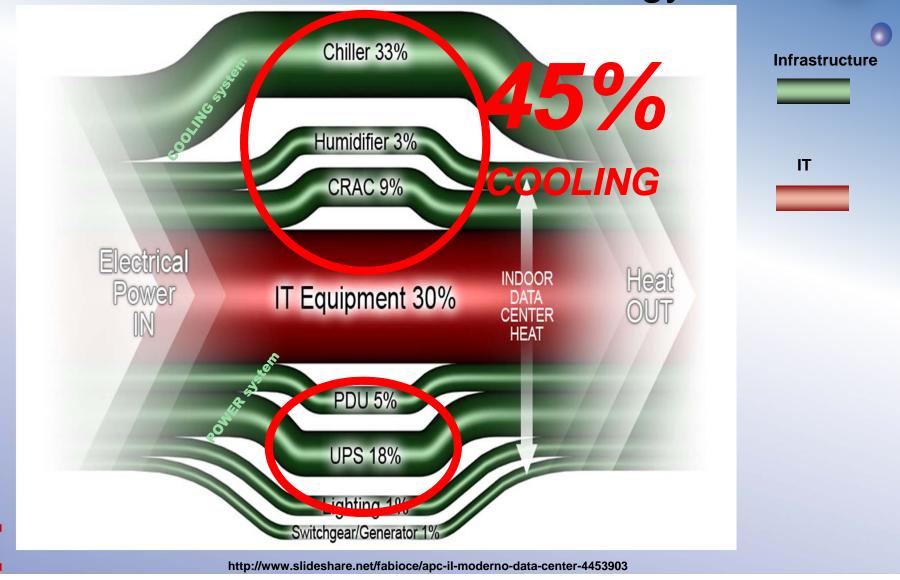






ADVANCED COMPUTING RESEARCH LABORATOR

Traditional Data Center Energy Use

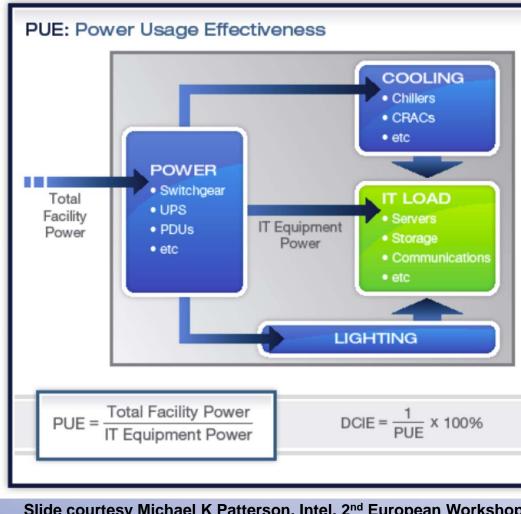






ADVANCED COMPUTING RESEARCH LABORATORY

Power Usage Effectiveness (PUE)



Slide courtesy Michael K Patterson, Intel, 2nd European Workshop on HPC Centre Infrastructure, Dourdan, France, 2010-10-06--08





Google

CHEP 2012 May 20 – 25, 2012 Lennart Johnsson

Q1 2011

Quarterly energy-weighted average PUE: **1.13** TTM energy-weighted avg. PUE: **1.16** Individual facility minimum quarterly PUE: **1.09**, Data Center E

Individual facility minimum TTM PUE*: **1.11**, Data Center J

Individual facility maximum quarterly PUE: 1.22 Data Center C

Individual facility maximum TTM PUE*: **1.21**, Data Center C

* Only facilities with at least twelve months of operation are eligible for Individual Facility Trailing Twelve Month (TTM) PUE reporting

 $PUE = \frac{E_{US1} + E_{US2} + E_{TX} + E_{HV}}{E_{US2} + E_{Net1} - E_{CRAC} - E_{UPS} - E_{LV}}$



- **EUS1** Energy consumption for type 1 unit substations feeding the cooling plant, lighting, and some network equipment
- **EUS2** Energy consumption for type 2 unit substations feeding servers, network, storage, and CRACs
- **ETX** Medium and high voltage transformer losses
- EHV High voltage cable losses
- ELV Low voltage cable losses
- ECRAC CRAC energy consumption
- **EUPS** Energy loss at UPSes which feed servers, network, and storage equipment
- **ENet1** Network room energy fed from type 1 unit substitution



http://www.google.com/corporate/datacenter/efficiency-measurements.html



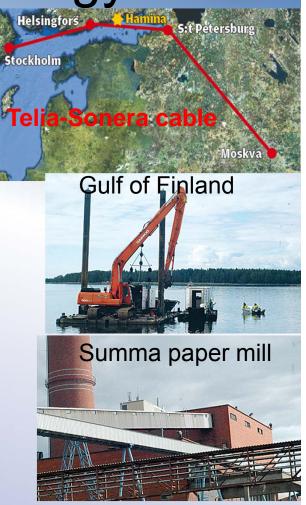


Google and Clean Energy

- The Hamina data center in Finland (previously the Summa paper mill)
 - Cooling water from Gulf of Finland (no chillers)
 - Four new wind turbines built
- Belgian data center designed without chillers. If the air at the Saint-Ghislain,

Belgium, data center gets too hot, Google shifts the data center's compute loads to other facilities.











ADVANCED COMPUTING RESEARCH LABORATOR

Facebook – Prineville Data Center







Facebook's Prineville, OR, 147,000-square-foot custom data center, with an estimated to cost \$188.2 million was brought into operation the summer of 2011. The site was chosen because of it's very dry and relatively cool climate. For 60 – 70% of the time cooling will be achieved by using cold air from outside. Excess heat from servers will be used to warm office space in the facility. **PUE 1.07 – 1.08**













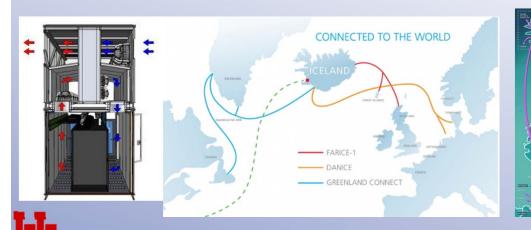
- Located at the Thor Data Center, Reykjavik
- Iceland Electric Energy 70% Hydro, 30% Geo Carbon Free, Sustainable

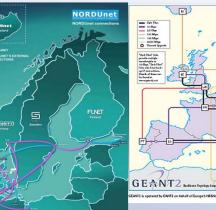


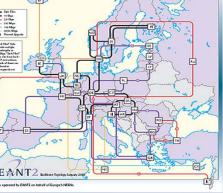




 Free Cooling – PUE in the 1.1 – 1.2 range; 1.07 for containerized equip. All time high temperature in Reykjavik: 24.8 C, Annual average ~5 C.





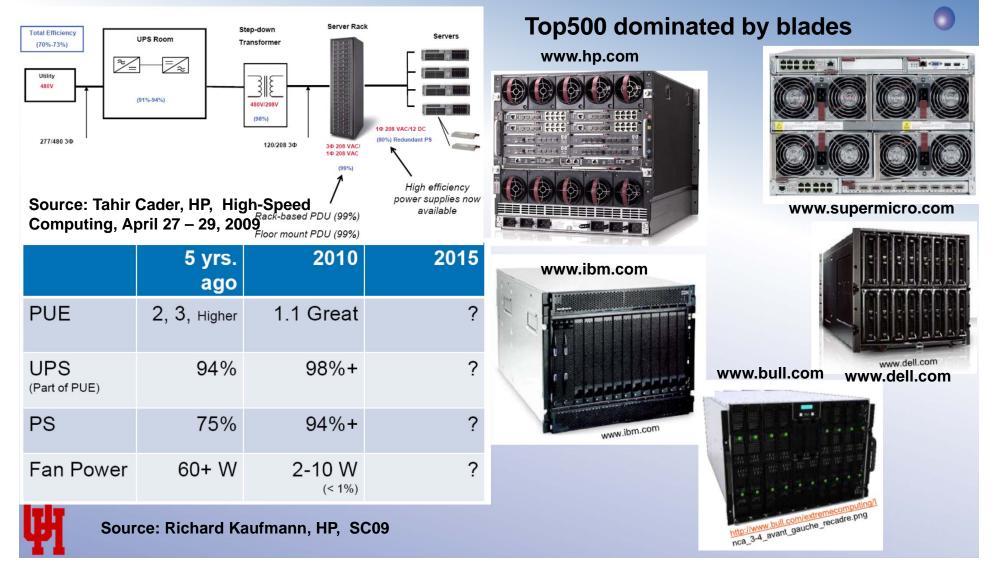






ADVANCED COMPUTING RESEARCH LABORATOR

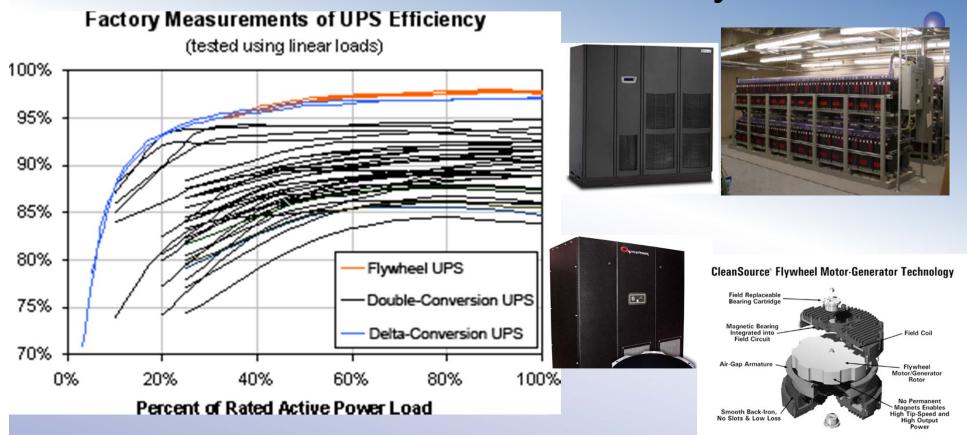
Data Center Power Efficiencies





ADVANCED COMPUTING RESEARCH LABORATOR

Data Center Power Efficiency - UPS



M. Ton, B. Fortenbury. December 2005. High Performance Buildings: Data Centers, Uninterruptible Power Supplies (UPS). LBNL. Ecos Consulting. EPRI Solutions. http://hightech.lbl.gov/documents/UPS/Final_UPS_Report.pdf





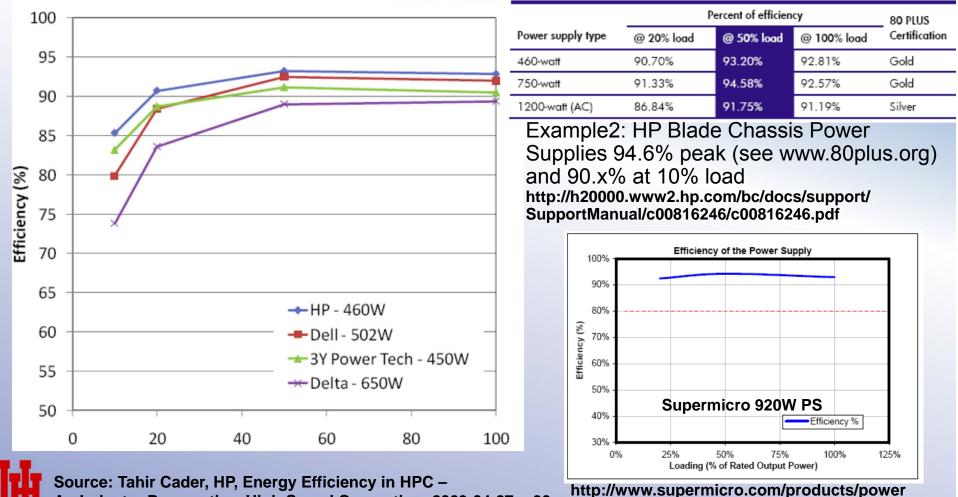


ADVANCED COMPUTING RESEARCH LABORATORY

Data Center Power Efficiency - PSU

Example1: HP Proliant Power Supplies

supply/80PLUS/80PLUS PWS-920P-1R.pdf



An Industry Perspective, High Speed Computing, 2009-04-27 -- 30



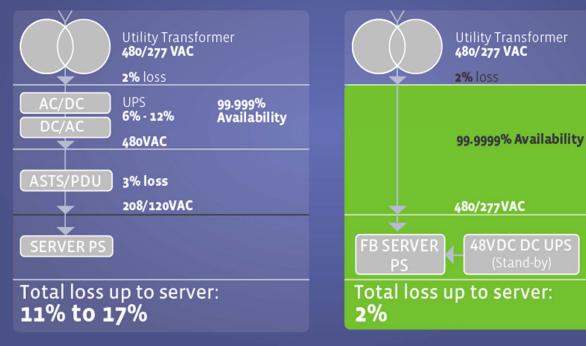


ADVANCED COMPUTING RESEARCH LABORATORY

Facebook – Prineville Data Center

Typical Data Center Power

Prineville Data Center Power



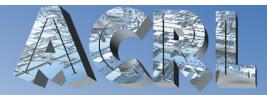
Google: UPS integrated with server PSU. UPS efficiency 99.9%

US Patent Office Application. June 1, 2007. Data Center Uninterruptible Power Distribution Architecture. <u>http://appft1.uspto.gov/netacg</u> <u>i/nph-</u> Parser?Sect1=PTO1&Sect2=

Parser?Sect1=P101&Sect2= HITOFF&d=PG01&p=1&u=% 2Fnetahtml%2FPTO%2Fsrch num.html&r=1&f=G&l=50&s1 =%2220080030078%22.PGN R.&OS=DN/20080030078&R S=DN/20080030078

Source: Amir Micahel, Facebook, August 17, 2011, http://www.hotchips.org/archives/hc23



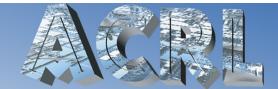


Data Center Energy Efficiency

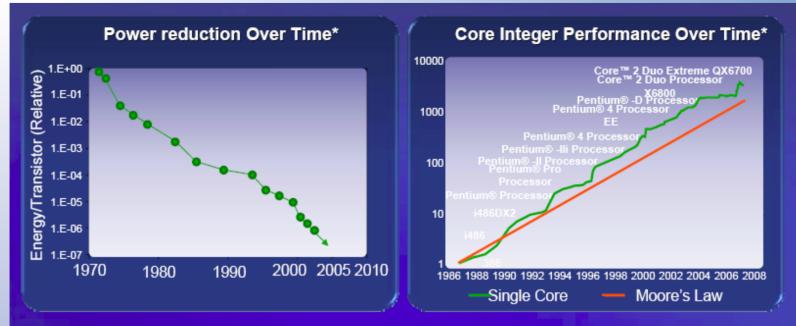
- Modern data centers are designed and operated for a PUE typically in the range of 1.05 – 1.2
- Future significant improvement in energy efficiency must come from architectures requiring less energy and applications that use them efficiently.







Incredible Improvement in Integrated Circuit Energy Efficiency



 ~ 1 Million Reduction In Energy/Transistor Over 30+ Years Delivering Great Performance Within Power Envelope Compute Energy Efficiency → Positive Impact On Environment



a

Source: Intel Corporate Technology Group

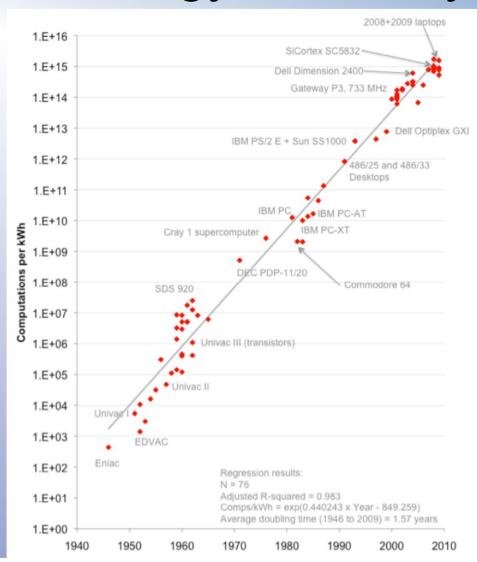
Source: Lorie Wigle, Intel, http://piee.stanford.edu/cgi-bin/docs/behavior/becc/2008/presentations/ 18-4C-01-Eco-Technology - Delivering Efficiency and Innovation.pdf





ADVANCED COMPUTING RESEARCH LABORATORY

Energy efficiency evolution



Energy efficiency doubling every 18.84 months on average measured as computation/kWh

Source: Assessing in the Trends in the Electrical Efficiency of Computation over Time, J.G. Koomey, S. Berard, M. Sanchez, H. Wong, Intel, August 17, 2009, http://download.intel.com/pressroom/pdf/comp utertrendsrelease.pdf

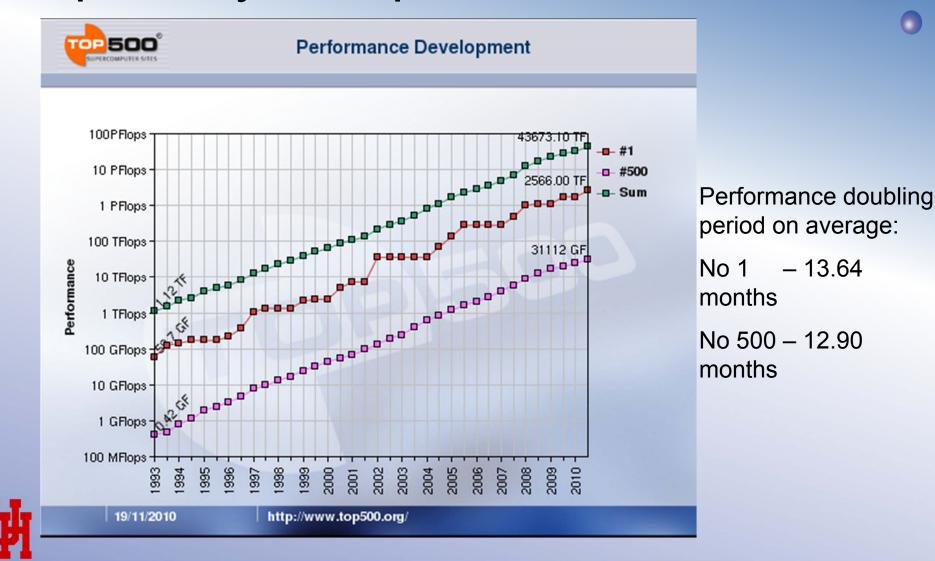






ADVANCED COMPUTING RESEARCH LABORATORY

Top500 system performance evolution





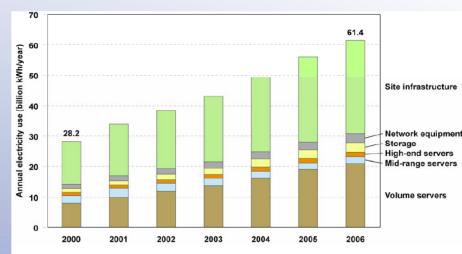


ADVANCED COMPUTING RESEARCH LABORATOR

The Gap

The energy efficiency improvement as determined by Koomey does not match the performance growth of HPC systems as measured by the Top500 list

The Gap indicates a growth rate in energy consumption for HPC systems of about 20%/yr.



	2000		2006		2000 - 2006
End use component	Electricity use	%	Electricity use	%	electricity use
	(billion kWh)	Total	(billion kWh)	Total	CAGR
Site infrastructure	14.1	50%	30.7	50%	14%
Network equipment	1.4	5%	3.0	5%	14%
Storage	1.1	4%	3.2	5%	20%
High-end servers	1.1	4%	1.5	2%	5%
Mid-range servers	2.5	9%	2.2	4%	-2%
Volume servers	8.0	29%	20.9	34%	17%
Total	28.2		61.4		14%

EPA study projections: 14% - 17%/yr Uptime Institute projections: 20%/yr PDC experience: 20%/yr

Report to Congress on Server and Data Center Energy Efficiency", Public Law 109-431, U.S Environmental Protection Agency, Energy Star Program, August 2, 2007, http://www.energystar.gov/ia/partners/prod_development/dow nloads/EPA_Datacenter_Report_Congress_Final1.pdf

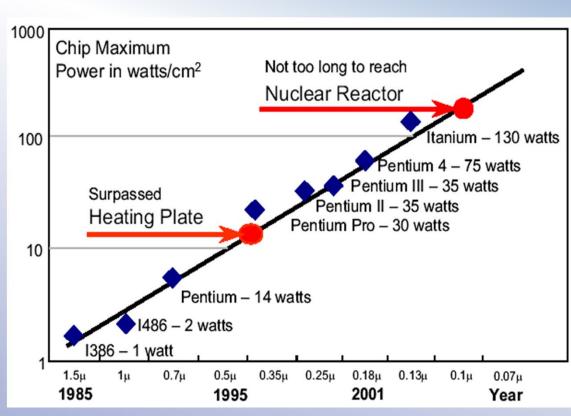
"Findings on Data Center Energy Consumption Growth May Already Exceed EPA's Prediction Through 2010!", K. G. Brill, The Uptime Institute, 2008, http://uptimeinstitute.org/content/view/155/147



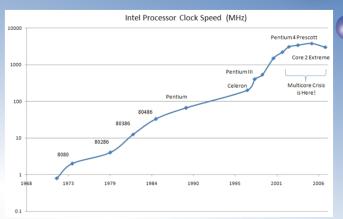


ADVANCED COMPUTING RESEARCH LABORATOR

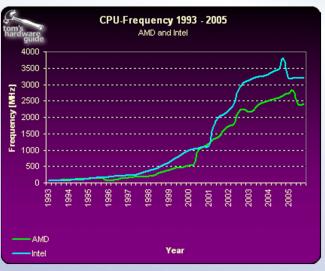
CPUs got hotter



Heat density of Intel CPUs, Source Shekhar Borkar, Intel



Intel Processor Clock Speed (MHz)), from http://smoothspan.files.wordpress.com/2007/09/clockspeeds.jpg



http://www.tomshardware.com/reviews/ mother-cpu-charts-2005,1175.html

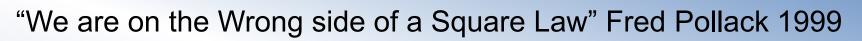






ADVANCED COMPUTING RESEARCH LABORATORY

Energy Consumption



New goal for CPU design: "Double Valued Performance every 18 months, at the same power level", Fred Pollack

Pollack, F (1999). *New Microarchitecture Challenges in the Coming Generations of CMOS Process Technologies.* Paper presented at the Proceedings of the 32nd Annual IEEE/ACM International Symposium on Microarchitecture, Haifa, Israel.

Ed Grochowski, Murali Annavaram Energy per Instruction Trends in Intel® Microprocessors. http://support.intel.co.jp/pressroom/kits/core2d uo/pdf/epi-trends-final2.pdf

Product	Normalized Performance	Normalized Power	EPI on 65 nm at 1.33 volts (nJ)
i486	1.0	1.0	10
Pentium	2.0	2.7	14
Pentium Pro	3.6	9	24
Pentium 4 (Willamette)	6.0	23	38
Pentium 4 (Cedarmill)	7.9	38	48
Pentium M (Dothan)	5.4	7	15
Core Duo (Yonah)	7.7	8	11

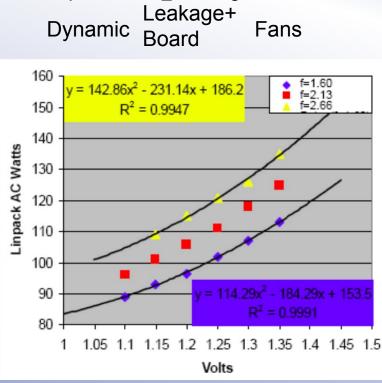




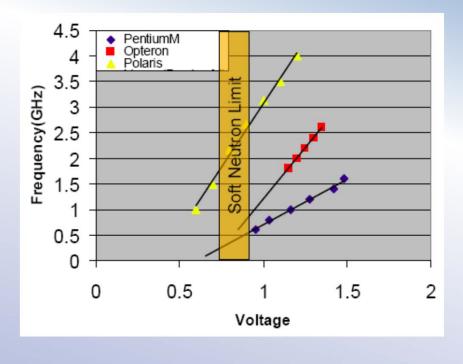
CED COMPUTING RESEARC

The Square Law

For CMOS the relationship between power (P), voltage (V and frequency (f) is $P = c_1 V^2 f + c_2 V + c_3 + O(V^4)$



Linpack: 15f(V-0.2)²+45V+19 STREAM: 5f(V-0.2)²+50V+19 Furthermore, $f \sim C(V-V0)$



Source: Supermicro

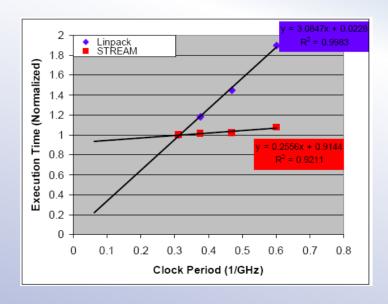




ADVANCED COMPUTING RESEARCH LABORATOR

The Square Law

• Execution Time $T=\alpha(1/f) + \beta$



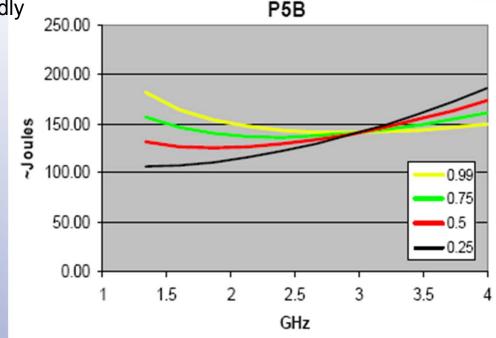
Putting it together: *Bus bound (black)*

ource: Supermicro

 minimize CPU power: lowest frequency

Putting it together:

Compute bound (yellow) – at low frequency the CPU leakage power and board combined with longer execution time increases the energy consumption; at high frequency dynamic and fan power increases rapidly P5B



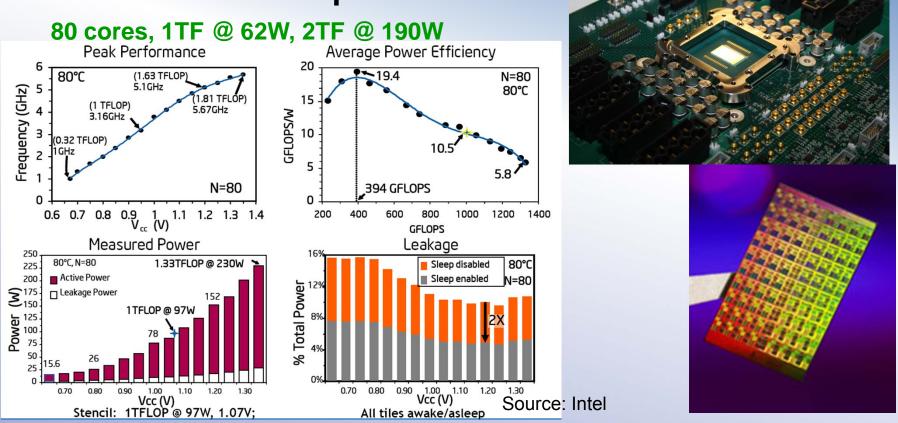




Power – Frequency Dependence

ICED COMPUTING RESEARCH LABORATOR

Intel Polaris Chip



S. Vangal, J. Howard, G. Ruhl, S.Dighe, H. Wilson, J. Tschanz, D. Finan, P. Iyer, A. Singh, T. Jacob, S. Jain, S. Venkataraman, Y. Hoskote, N. Borkar. February 11-15, 2007. An 80-tile 1.28 Tflops Network-on-Chip in 65 nm CMOS. Pp. 98 – 99. IEEE Solid-States Circuits Conference, San Francisco. http://ieeexplore.ieee.org/xpl/freeabs_all.isp?arnumber=4242283





Some Low Frequency designs

- Scicortex, MIPS 800 MHz
- Blue Gene/L (750 MHz), P (800 MHz) and Q (1.6 GHz)
- DSPs ~1 GHz (Ex. TI TMS320C6678)
- ARM < 1GHs 2 GHz (Server Ex. Calxeda HP Moonshot)
- Mobile low to high hundreds MHz
- Greenflash







Green Flash Strawman System Design

Three different approaches examined (in 2008 technology) Computation .015°X.02°X100L: 10 PFlops sustained, ~200 PFlops peak

- AMD Opteron: Commodity approach, lower efficiency for scientific applications offset by cost efficiencies of mass market
- BlueGene: Generic embedded processor core and customize systemon-chip (SoC) to improve power efficiency for scientific applications
- Tensilica XTensa: Customized embedded CPU w/SoC provides further power efficiency benefits but maintains programmability

Processor	Clock	Peak/ Core (Gflops)	Cores/ Socket	Sockets	Cores	Power	Cost 2008			
AMD Opteron	2.8GHz	5.6	2	890K	1.7M	179 MW	\$1B+			
IBM BG/P	850MHz	3.4	4	740K	3.0M	20 MW	\$1B+			
Green Flash / Tensilica XTensa	650MHz	2.7	32	120K	4.0M	3 MW	\$75M			
Slide courtesy Horst Simon, NERSC, http://www.cs.berkeley.edu/~demmel/cs267_Spr09/Lectures/SimonPrinceton0409.ppt										







DARPA Exascale study

- Last 30 years:
 - "Gigascale" computing first in a single vector processor
 - "Terascale" computing first via several thousand microprocessors
 - "Petascale" computing first via several hundred thousand cores
- Commercial technology: to date
 - Always shrunk prior "XXX" scale to smaller form factor
 - Shrink, with speedup, enabled next "XXX" scale
- Space/Embedded computing has lagged far behind
 - Environment forced implementation constraints
 - Power budget limited both clock rate & parallelism
- "Exascale" now on horizon
 - But beginning to suffer similar constraints as space
 - And technologies to tackle exa challenges very relevant

屮

Especially Energy/Power

http://www.ll.mit.edu/HPEC/agendas/proc09/Day1/S1_0955_Kogge_presentation.ppt





Power fundamentals – Exascale

Processor

- Modern processors being designed today (for 2010) dissipate about 200 pJ/op total. This is ~200W/TF 2010
- In 2018 we might be able to drop this to 10 pJ/op
 - ~ 10W/TF 2018
- This is then 16 MW for a sustained HPL Exaflops
- This does not include memory, interconnect, I/O, power delivery, cooling or anything else

Memory

- Cannot afford separate DRAM in an Exa-ops machine!
- Propose a MIP machine with Aggressive voltage scaling on 8nm
- Might get to 40 KW/PF –

60 MW for sustained Exa-ops



Source: William J Camp, Intel, http://www.lanl.gov/orgs/hpc/salishan/pdfs/Salishan%20slides/Camp2.pdf





Power fundamentals - Exascale

Interconnect

- For short distances: still Cu
- Off Board: Si photonics
- Need ~ 0.1 B/Flop Interconnect
- Assume (a miracle)
 5 mW/Gbit/sec
 - ~ 50 MW for the interconnect!

Power and Cooling

Still 30% of the total power budget in 2018! Total power requirement in **2018**: **120—200 MW!**



Source: William J Camp, Intel, http://www.lanl.gov/orgs/hpc/salishan/pdfs/Salishan%20slides/Camp2.pdf

1/0

- Optics is the only choice:
- 10-20 PetaBytes/sec
- ~ a few MW (a swag)





ADVANCED COMPUTING RESEARCH LABORATOR

What type of Architecture?



Reducing Waste

Mark Horowitz 2007: "Years of research in lowpower embedded computing have shown only one design technique to reduce power: <u>reduce waste</u>."

Seymour Cray 1977: "Don't put anything in to a supercomputer that isn't necessary."







Exascale Computing Technology Challenges, John Shalf National Energy Research Supercomputing Center, Lawrence Berkeley National Laboratory ScicomP / SP-XXL 16, San Francisco, May 12, 2010





ADVANCED COMPUTING RESEARCH LABORATOR

What type of Architecture?

Instruction Set Architecture



A Short List of x86 Opcodes that Science Applications Don't Need!

menonic	<u>op1</u>	<u>sp2</u>	<u>op2</u>	op.4	Seitt	22 5	rpe	KI 0	pisc	-	n II	2	asted f	h liber	4+1 1	under f	f valuer	description, notes
44	22.	231				П	37					-		0		1 m.p.		#3011 Adgust Attes Addstinn
.sD	12	an					0.8	40						oemape		1		\$3011 Adjust Al Batone Division
181	<u>115.</u>	au					0.9	0.K						0 snap c	····##.p.	1		ASCII Adjust AT After Hultiply
88	55	AH					37						· · · · · ·	2		1		63011 Adjust AL After Subtraction
ac	z/m8	2.8				\square	20	x				1.		0	0			6dd with Carry
ae	c/m16/32/64	13/01/69				T	11	x				L.		0	0			2dd with Carry
ae.		afeelt		1		Ħ	32					-			*******			Add with Carey
ne	-46/10/64	a/ral#/02/54		+		Ħ	12					1		*···*****	********			Add with Carey
ae	ar.	inerall		+		Ħ	24											Add with Carey
DC	x8X	iamL6/32	l	+		++	23	+	-		++			a	0	-		Add with Carry
ac	x/340	ianó		+	-	++	0.0	x	-	-	++	1		aapc	0	+		Add oith Carry
DC.	s/a15/32/64	lanL6/32		+	-	++	01	2	-	++	+ +	1		a	0KE20C	+		Add with Carry
DC .	12/240	Larvô		+	-	++	20	2	-	-	-	_		0	0FEADC	+	-	add with Carry
20	w/a18/32/64	Land		+	-	++	0.0	2	-		_	_		0	*	-		345 with Carry
up	2/28	25	-	-		++	00		-		-	1		0	07840C	-		add
0.0	2/m2 2/m18/32/64	210/02/69	-	-	-	++	00		-			1		0	07840C	-		845
20	2/311/32/64	2/08		-		++	02	1	-		+	~		0	07840C	-		845
	v18/32/84	2/10/02/14	l	+	-	++	-		-		+ +	-				-		848
20				-	-	++	02		-		+ +	-		0	0	-		
20	8L	Lars3		-	-	++	0.3	-	-	\vdash	+ +			0 98 Apr	0954pc	-		840
0.0	cfiX	SC/Mierei		-	-	++	0.0	-	-	\vdash				05x4pt	0	-		840
0D	c/a8	iard .		-		\square	80	0				L		0	095ADC	-		ðdð
an	c/a16/32/64	isersL6/02					18	0				L		07¥#pt	055ADC			રેનેતે
an	n/m2	ined			_		92	a.				1		* • • • • • •	*******	-		242
an	=/m16/22/64	ined	-				92	Ø,				1		*···*****	********			242
CIDP3	885	22014,3120			****		T 50		241									Add Packed Double-TP Values
CIDP 1	X805	2201(1120			*xel		F 56		P21									Add Packed Single-TP Values
20053	3335	ann/ad4			exel	72 0	r 50	x	24+									Add Scalar Double-TP Values
2053	385	sane's 32			scel	72 0	F 56	x	P2+									Add Scalar Single-TP Values
CECTORS	X805	apre/ # 120			arel.	55 0	r Do	x	2-9++									Packed Double-TP Add/Dubyracy
C SCITCOL	3325	anny's 120		-	arel.	72 0	r Do	×	2-4++									Packed Single-TV Add/Duburacu
ex.	105	231	ira 5	-		+	0.5							0		1		Sdjurt &Z Bafore Division
S87.0				-		69			P4+	այ								Alternating branch prefix (used only with dec instructions)
152	80	an .	ine 5	+	-	-	0.9	+		-	++			0		***** * .*		ödýurð ál Afber Bulbiply
an	2/38	10	11000	+	-	++	20		-	++	++	1		0	095.pc			Logical AND
80	c/a16/32/64	+16/02/69		1		++	21	1	-		++	1.		0	099.pc			Logical AND
an	=0	2.107 0 17 0 2		1		++	22	1			++	-		0	099.pc			Logical AD
an	+15/33/64	a/m14/32/64		-		++	29	-			++				R		100-000000000	Legisal AD
30	AL.	ianó		-	-	++	24		-		++			0	6			Logical MD
30 30	AL.			-	-	++	25	+	-	-	++	+				-		Logical AD Logical AD
30		ian16/32		-	-	++	_	-	-		++	1		0 102 pt	eec.pc	· · · · · 3 · ·		
	k/aŭ	iané		+	-	++	āé	4			_	-		a taape	0EE.pc	· · · · · · 3 · ·		Logical AND
370	x/a18/32/64	lan16/32		-			01	4	_			1		ataapt	6EE.pc			Logical AND
510	2/240	Lano					20	4				1		osmapc	0P8.pc			Logical AND
370	2/2418/32/64	Land					40		00+			l		0	0FE.pc	· · · · · · · · ·		Logical AND
80920	3035	35N/3128			2762				P-9+									Sitwire Legical AND NOT of Packed Double-IP Valuer
BDSD 3	3035	25m/2128			***1	-	r 85		50+									Birwire Logical AND BUT of Packed Single-IP Values
8093	X85	2214 ¹ 9128			***2	68 0	r 54		p.q+									Strwigs Logical AND of Packed Double-IP Values
SDP1	3005	2305 ¹ 2128			Teet	1	r 89	T	p.;+									Sitwire Logical ASD of Parked Single-EP Valuer



Exascale Computing Technology Challenges, John Shalf, *National Energy Research Supercomputing Center, Lawrence Berkelev National Laboratory* ScicomP / SP-XXL 16, San Francisco, May 12, 2010





ADVANCED COMPUTING RESEARCH LABO

What type of Architecture?



z16/01 s16/01116/02 1/m16/22/64 -18/55/81 wat 122/22 Um16732/64 13.6/327.64 2/10/32/61 r/m16/32/64 126/22/04 (m16/32/64 x1m16/32/64 r/m16/32/64 15(31/64 min.16/32/64 813102 2/104 p1116:16/0 146:16 *16/32 C100A26 CERVISIC #16/32 x10/37 CRONINA CROAL CRONNED CHINAR x16/32/ +16/32 CE01056 001033 CROWNE ¥\$6/32 +16/32 0010032 x16/32

More Wasted Opcodes

CU	150251	x32/64	×nm/ = 64						
cu	77 502 55	an	×100/264				FRCHA	52	STL
	7151250	ann.	2/222/64	r36/32	/01 z/mit/32/8				
51 CV	151253	ann.	z/w72/64	-56/32			- TWCR4	sr	STL
	T 552.5D	an.	sing/se7z	£16/32			EWCH?	52	STL
	155251	252/64	11NN/317Z	*/**	•8		236027	52	STL.
CV	/TTFD2DQ	arn.	nnav/ar128	2/11/2/	22/54 z15/22/54		FERSTOR	527	527.2
CU:	TTPD2PI	nn.	nnav/ar128	53	2/05		TIRSTOR	57	371
r cu	TTPB2D0	8 n	nnes/ss128	130 /32.3	/64 z/mLE/32/64	9	FISAVE	n512	57
2 00	TTPBEPT	-	nma/ a64	77	inerali	-			
	TTSDAST	x32/64	×na/264	283	1an15/22	_	FESAUE	n512	52
-	7755231	x32/64	xmm/m22	5(a)D	Larvó		ECTBACT	28	
- 1				c/si36/	22/84 tanL5/02		FALSI	881	87
cu		13.	AX .	s/#8	Lars5				
00		13.	JOY .	*/*14/	22/54 invit		- FYLZIPL	527	57
cn,	10	FRY	2.AX	3080	2004 0120	iané	G8	65	
00	10	RDX	8.48	2000	2007e/30100	Laund	HADDED	1071030	xnay n126
0 00	TO:	ERY	87	51	r 8		HADDPS	8(9)7	nnav n120
7 D.A	u	AL.		m/	nd.	_	- MLT		
17 D.A.	u	AL		316	r3.0		11303.20	ana.	2000/0120

•We only need 80 out of the nearly 300 ASM instructions in the x86 instruction set!

- Still have all of the 8087 and 8088 instructions!
- Wide SIMD Doesn't Make Sense with Small Cores
- Neither does Cache Coherence
- •Meither does HW Divide or Sqrt for loops

CUTPBEPD

- Creates pipeline bubbles
- Better to unroll it across the loops (like IBM MASS libraries)
- •Move TLB to memory interface because its still too huge (but still get #16/32 *16/32

precise exceptions from segmented protection on each core)

#16/32/ 216/32 CROWE Science U.S. DEPARTMENT OF ENERGY

*16/32)

C2077A CRIMER,

CHINGE

Exascale Computing Technology Challenges, John Shalf National Energy Research Supercomputing Center, Lawrence Berkeley National Laboratory ScicomP / SP-XXL 16, San Francisco, May 12, 2010

13770

INVLPO

#Tlags



What kind of architecture (core)

Xtensa x 3		•
Tensilica	DP ARM	•
Intel Core2		•
Po	wer 5	•
L3 directory	sontrol	•

ERSC

How Small is Small

- Power5 (server)
 - 389mm^2
 - 120W@1900MHz
- Intel Core2 sc (laptop)
 - 130mm^2
 - 15W@1000MHz
- ARM Cortex A8 (toaster oven)
 - 5mm^2
 - 0.8W@800MHz
- Tensilica DP (cell phones)
 - 0.8mm^2
 - 0.09W@600MHz
- Tensilica Xtensa (Cisco Rtr)
 - 0.32mm^2 for 3!
 - 0.05W@600MHz

Cubic power improvement with lower clock rate due to V²F

Slower clock rates enable use of simpler cores

Simpler cores use less area (lower leakage) and reduce cost

Tailor design to application to <u>reduce</u> <u>waste</u>

Each core operates at 1/3 to 1/10th efficiency of largest chip, but you Science can pack 100x more cores onto a chip and consume 1/20 the power

http://www.csm.ornl.gov/workshops/SOS11/presentations/j_shalf.pdf





ADVANCED COMPUTING RESEARCH LABORATOR

What kind of architecture - accelerators

				CBE	QS20 Blade
	Cell BE	Nvidia G80 GF100	ClearSpeed CSX600		
32-bit FP	200+ GFLOPS	360+ GFLOPS	25+ GFLOPS		
64-bit FP	20+ GFLOPS 200+G		25+ GFLOPS 96 (GF	
Clock frequency	3.2 GHz	575 MHz	210 MHz		
Transistors/ chip	~ 241M	~ 681M	~ 128M	Clearspeed 1U 0.97TF	
Power	~ 110 Watts	~ 145 W (for 225W GeForce 8800 GTX board)	25W board		
GF8800GTX	Т	/idia	eessor		CBE Blade
http:/	//gamma.cs.unc.edu/SC	ClearSp	hip Bridge Ports eed CSX600 ulticore-Workshop.	Clearspeed PCI-X board	

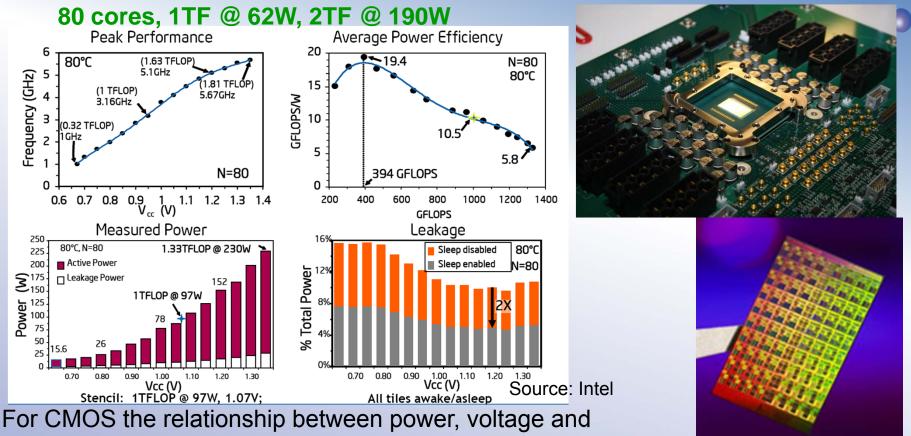




ADVANCED COMPUTING RESEARCH LABO

CHEP 2012 May 20 – 25, 2012 Lennart Johnsson

What type of architecture?



frequency is

$P = c1V^2f + c2V + c3 + O(V^4)$

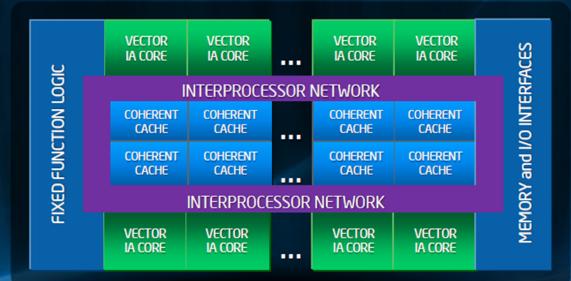
S. Vangal, J. Howard, G. Ruhl, S.Dighe, H. Wilson, J. Tschanz, D. Finan, P. Iyer, A. Singh, T. Jacob, S. Jain, S. Venkataraman, Y. Hoskote, N. Borkar. February 11-15, 2007. An 80-tile 1.28 Tflops Network-on-Chip in 65 nm CMOS. Pp. 98 – 99. IEEE Solid-States Circuits Conference, San Francisco. http://ieeexplore.ieee.org/xpl/freeabs all.isp?arnumber=4242283



ADVANCED COMPUTING RESEARCH LABORATOR

CHEP 2012 May 20 – 25, 2012 Lennart Johnsson

Intel[®] MIC Architecture: An Intel Co-Processor Architecture



Many cores and many, many more threads Standard IA programming and memory model







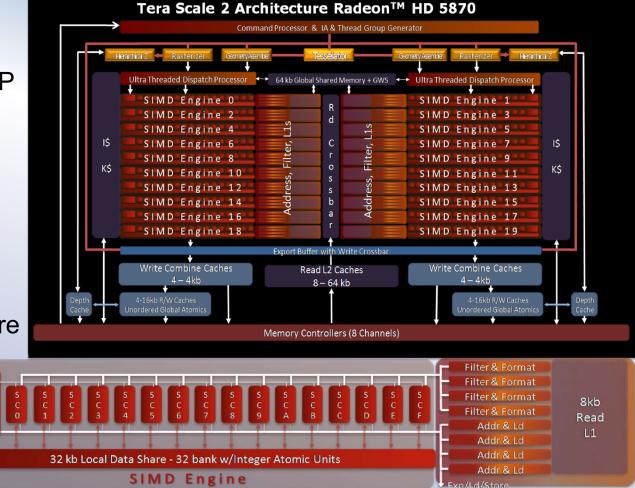
GPUs – AMD 5870 (2010)

- 1600 PEs
- 20 SIMD Engines (SE)
- 2.72 TF SP, 0.544 TF DP
- Memory BW 147GB/s
- 8kB L1 and 32kB data share for each SE
- 64kB Global data share
- Four 128 kB L2 caches
- Up to 272 billion 32-bit fetches/second
- Up to 1 TB/sec L1 texture fetch bandwidth

(Branch

Unit)

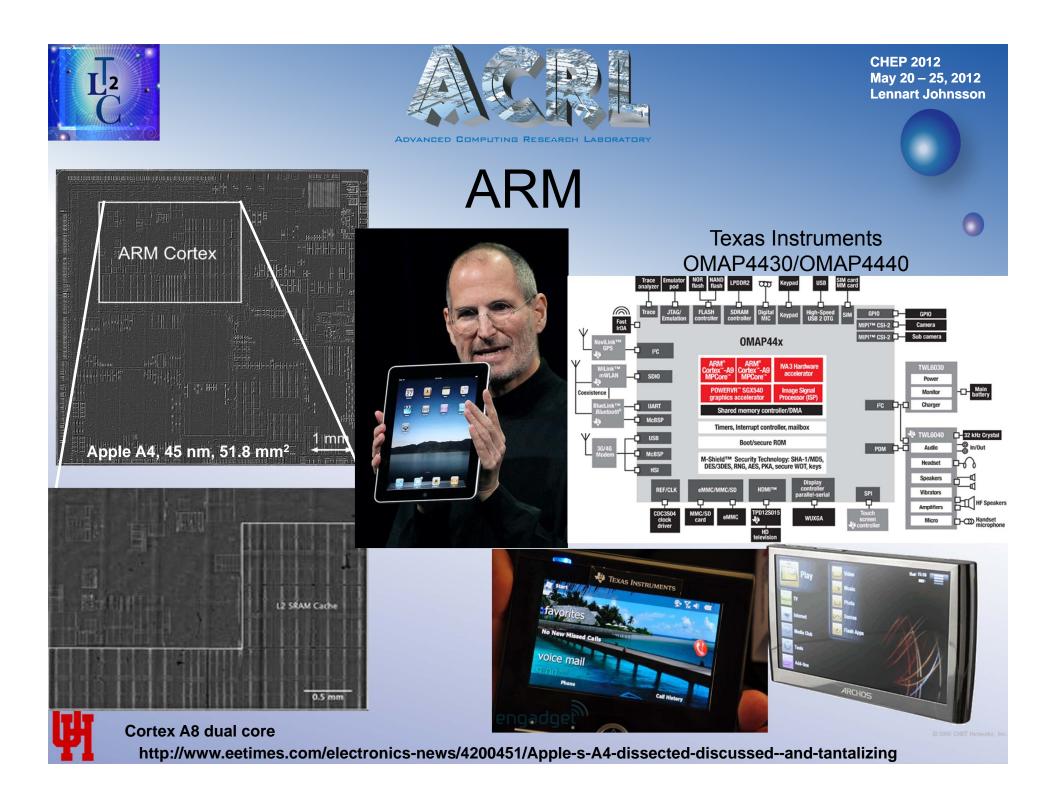
- Up to 435 GB/sec between L1 & L2
- 225W







	Advanced Computing Research Laboratory							
Green500 Rank June 2011	MFLOPS/W	Site*	Computer*	Total Power (kW)				
<u>1</u>	2097.2	IBM Thomas J. Watson Research Center	NNSA/SC Blue Gene/Q Prototype 2	40.95				
<u>2</u>	1684.2	IBM Thomas J. Watson Research Center	NNSA/SC Blue Gene/Q Prototype 1	38.8				
<u>3</u>	1375.9	Nagasaki University	DEGIMA Cluster, Intel i5, ATI Radeon GPU, Infiniband QDR	34.24				
<u>4</u>	958.35	GSIC Center, Tokyo Institute of Technology	HP ProLiant SL390s G7 Xeon 6C X5670, Nvidia GPU, Linux/Windows	1243.8				
<u>5</u>	891.88	CINECA / SCS - SuperComputing Solution	iDataPlex DX360M3, Xeon 2.4, nVidia GPU, Infiniband	160				
<u>6</u>	824.56	RIKEN Advanced Institute for Computational Science (AICS)	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect	9898.6				
<u>7</u>	773.38	Forschungszentrum Juelich (FZJ)	QPACE SFB TR Cluster, PowerXCell 8i, 3.2 GHz, 3D- Torus	57.54				
<u>8</u>	773.38	Universitaet Regensburg	QPACE SFB TR Cluster, PowerXCell 8i, 3.2 GHz, 3D- Torus	57.54				
<u>9</u>	773.38	Universitaet Wuppertal	QPACE SFB TR Cluster, PowerXCell 8i, 3.2 GHz, 3D- Torus	57.54				
<u>10</u>	718.13	Universitaet Frankfurt	Supermicro Cluster, QC Opteron 2.1 GHz, ATI Radeon GPU, Infiniband	416.78				
<u>11</u>	677.12	Georgia Institute of Technology	HP ProLiant SL390s G7 Xeon 6C X5660 2.8Ghz, nVidia Fermi, Infiniband QDR	94.4				
<u>12</u>	650.3	National Institute for Environmental Studies	Asterism ID318, Intel Xeon E5530, NVIDIA C2050, Infiniband	115.87				
<u>13</u>	635.15	National Supercomputing Center in Tianjin	NUDT TH MPP, X5670 2.93Ghz 6C, NVIDIA GPU, FT- 1000 8C	4040				
<u>14</u>	565.97	Yukawa Institute for Theoretical Physics (YITP)	Hitachi SR16000 Model XM1/108, Power7 3.3Ghz, Infiniband	129.6				
<u>15</u>	555.5	CSIRO	Supermicro Xeon Cluster, E5462 2.8 Ghz, Nvidia Tesla s2050 GPU, Infiniband	94.6				
<u>16</u>	492.64	National Supercomputing Centre in Shenzhen (NSCS)	Dawning TC3600 Blade, Intel X5650, NVidia Tesla C2050 GPU	2580				
<u>17</u>	483.66	IBM Thomas J. Watson Research Center	Power 750, Power7 3.86 GHz, 10GigE	120.56				
<u>18</u>	467.73	CeSViMa - Centro de Supercomputación y Visualización de Madrid	BladeCenter PS702 Express, Power7 3.3GHz, Infiniband	154				







Start-Up Aims to Slay Chip Goliath

By ASHLEE VANCE Published: August 15, 2010

A group of investors, including companies from the United States, Europe and the United Arab Emirates, has formed in a bid to disrupt one of <u>Intel</u>'s most lucrative franchises.

🕀 Enlarge This Image



Ben Sklar for The New York Times Barry Evans is chief of Smooth-Stone, a name that refers to David's weapon in the Bible.

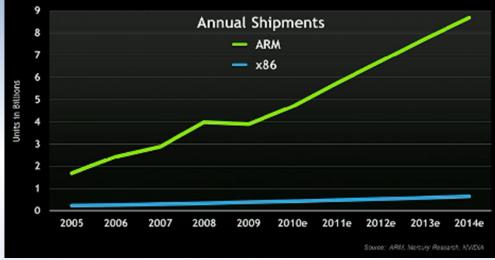
The companies have put \$48 million into Smooth-Stone, a start-up based in Austin, Tex., betting that it can modify low-power smartphone chips to run servers, the computers in corporate data centers. If successful, Smooth-Stone would undermine Intel's server-chip business and offer companies, especially those with vast data centers like <u>Google</u>,

Amazon.com, Facebook and Microsoft,

cost savings.

中

ARM is Pervasive and Open



January 06, 2011 NVIDIA ARMs Itself for Heterogeneous Computing Future

..... On Wednesday, the GPU-maker -- and soon to be CPU-maker -- revealed its plans to build heterogeneous processors, which will encompass high performance ARM CPU cores alongside GPU cores. The strategy parallel's AMD's Fusion architectural approach that marries x86 CPUs with ATI GPUs on-chip.





ADVANCED COMPUTING RESEARCH LABORATORY

Digital Signal Processors and HPC?

Texas Instruments 8-core DSP TMS320C6678

- Industry-best floating point performance
 - 16 Gflops/W
- Standard programming model
 supports MPI and OpenMP
- Wide range of applications
 - from embedded systems to server blades
- Full ecosystem support
 - Off the shelf PCIe and ATCA cards
 - O/S and application software

Supported by a full set of development tools and Code Composer Studio IDE

MS320C667









ADVANCED COMPUTING RESEARCH LABORATORY

Nominal Energy Efficiency of Mobile CPUs, x86 CPUs and GPUs

ARM Cortex-9		tex-9	۵	TON	n	AM	0 12-0	core	Inte	el 6-c	ore	ATI 9370		
Cores	W	GF/W	Cores	W	GF/W	Cores	W	GF/W	Cores	W	GF/W	Cores	W	GF/W
4	~2	~0.5	2	2+	~0.5	12	115	~0.9	6	130	~0.6	1600	225	~2.3

nVidia Fermi TMS320C6678							3M BQ	С	ClearSpeed CX700			
Cores	W	GFJW	Cores	W	GF/W	Cores	W	GF/W	Cores	W	GF/W	
512	225	~2.2	8	4	~ 15	16	55	3.7	192	10	~10	

Very approximate estimates!!



KTH/SNIC/PRACE Prototype II





ADVANCED COMPUTING RESEARCH LABORATOR

A DSP Example (2011): A Voice and Video Processing Board

- 1.2 TF Peak Double Precision (DP) (3.2 TF Peak Single Precision (SP))
- 20 GB memory
- 256 GB/s memory bandwidth
- 240W
- 2 nano Joules/DFLOP (5 GF/W, DP)
- 100 Gbps interconnect total bandwidth
- Dual 10Gbps Ethernet uplink
- 20 devices, 8 cores each
- 50 Gbps links pairing devices



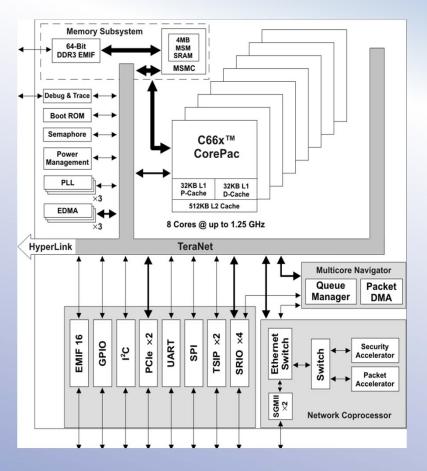


Source: Pekka Varis, TI





Texas Instruments TMS320C6678



मि

Image © Texas Instruments

8 C66x Cores:

4+4 wide VLIW, in-order, A/B-side 4 DP-add, 2 DP-mul, 2 load/store 32+32 registers, 32-bit wide

Memory System:

32 kB L1 data and program SRAM
512 kB L2 unified SRAM
4 MB shared L3 SRAM - MCSM
8 GB DDR3-1600 (1333), 64 bit wide

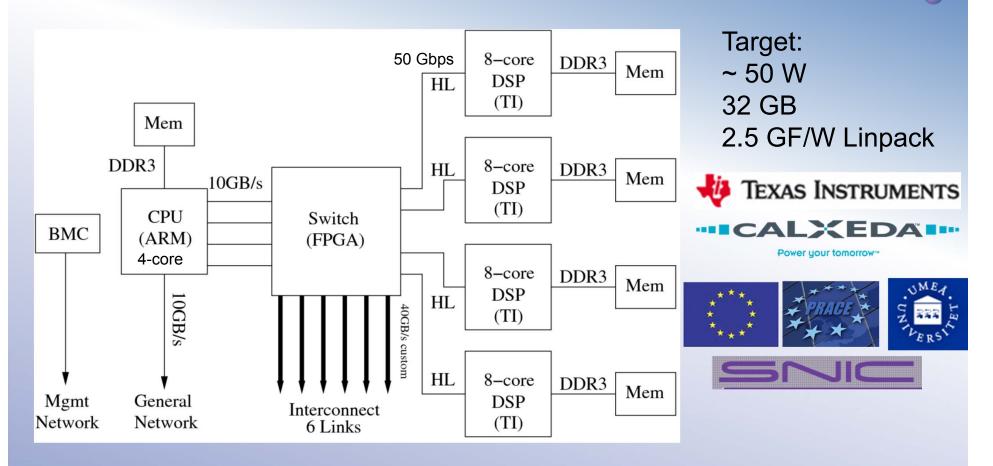
Communication:

- 2 gigabit Ethernet ports
- 1 serial rapid IO 4x5 Gbps 57
- 1 HyperLink 4x12.5(10) Gbps



ADVANCED COMPUTING RESEARCH LABORATORY

KTH/SNIC/PRACE DSP HPC node

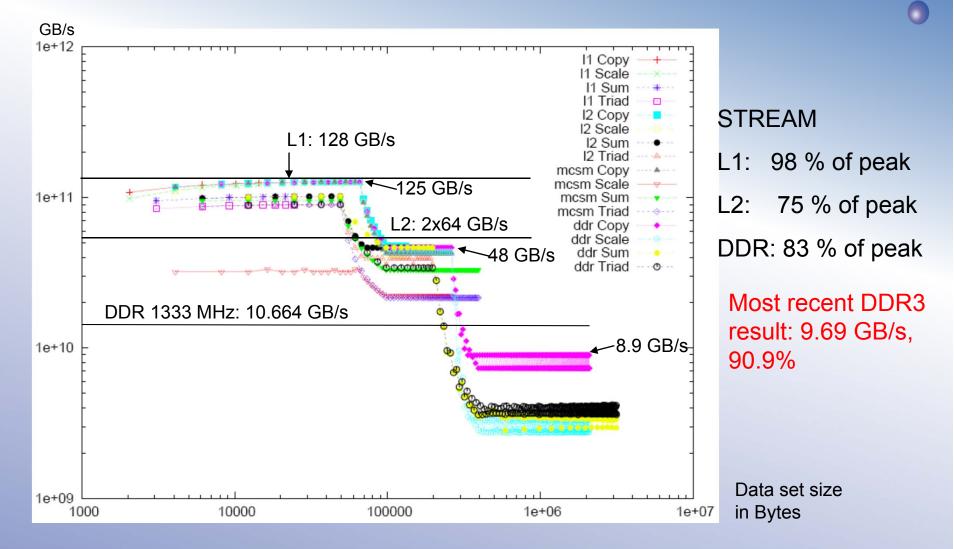






ADVANCED COMPUTING RESEARCH LABORATORY

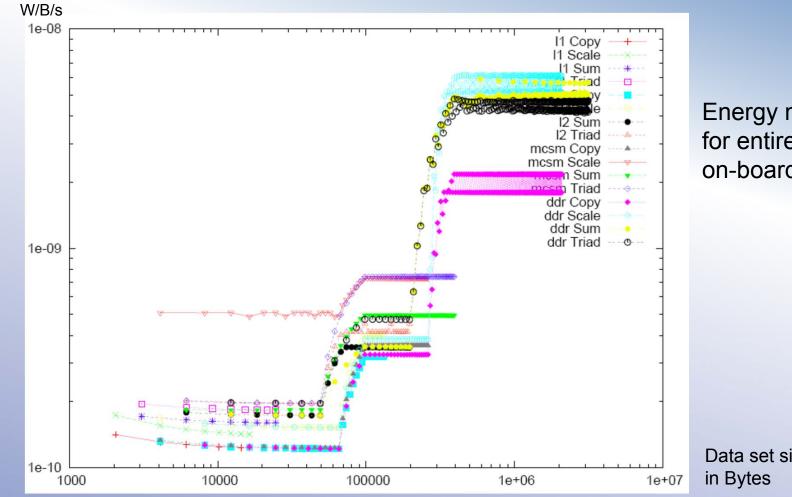
STREAM 6678 Bandwidth test 8 cores





ADVANCED COMPUTING RESEARCH LABOR

STREAM 6678 Bandwidth test 8 cores



Energy measured for entire EVM with on-board emulator

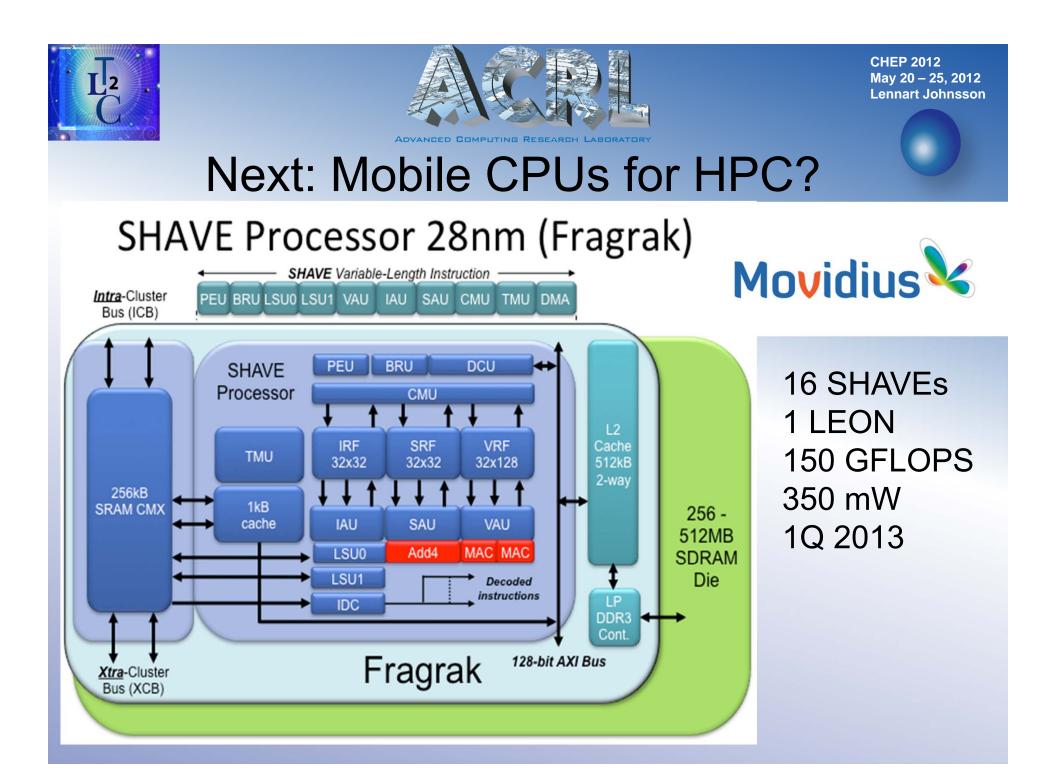
Data set size







- Innermost loop L1: >95%
- Current 8-core result: 49%, expected after further optimization >75% (comparable to Interlagos result reported by EPCC, but less than Westmere)
- Expected energy efficiency comparable to Blue Gene/Q







63

ADVANCED COMPUTING RESEARCH LABORATORY

SHAVE Performance

	SHAVE Fragrak	BlueGene/Q	NVIDIA Kepler
Clock Frequency	800 MHz	1600 MHz	1006 MHz
Cores/Threads	16	16	1536
FP Performance	51.2 GF/s	204.8 GF/s	~ 1000 GF/s
Power	0.35 W	55 W	195 W
Memory	512 MB	16 GB	2 GB
Memory Bandwidth	6.4 GB/s	42.7 GB/s	192.2 GB/s
Network Bandwidth	1 GB/s	22 GB/s	12.8 GB/s
Energy Efficiency	146 GFLOP/J	3.7 GFLOP/J	5.1 GFLOP/J
FLOP/Memory Cap	95 FLOP/B	12 FLOP/B	466 FLOP/B
FLOP/Memory BW	7.5 FLOP/B	2.5 FLOP/B	5.2 FLOP/B
FLOP/Network BW	48 FLOP/B	9.3 FLOP/B	78 FLOP/B

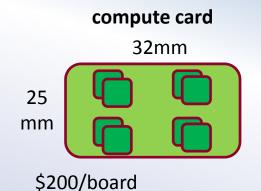






ADVANCED COMPUTING RESEARCH LABORATORY

Movidius 10 PFLOPS Strawman



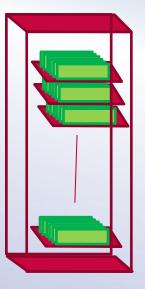
615 DP GFLOPS @ 2.8W

8*128MB DDR3 @ 1.2 GHz 76.8GB/s Mem BW (8*9.6)

(8* 4 * 16 * 800MHz)

Node card 9840 TFLOPS 16x compute cards 45W

Cabinet 10 petaFLOPS 1024 Nodes 46kW 40 sq ft



10 PFLOPS in a single cabinet





ADVANCED COMPUTING RESEARCH LABORATORY

10 PFLOPS Comparison

											Netwk						
			Flops/	Pk Core			Sub-		Mem	Pk	BW		Tot.		Tot.		
			Clock/	GFLOP	Cores/	Watts/	domains	MBytes	BW	Bytes/	(GB/s	# M	Power	Tot.	petaFLOP		
Name	CPU	GHz	Core	S	Socket	Sckt	/Sckt	/Socket	GB/s	FLOP)	Sockets	MW	Cost \$M	S	\$/socket	\$/GFLOP
AMD	Opteron	2.8	2	5.6	2	95	22.4	112	6.4	0.57	0.57	0.89	179	1799.6	9.97	2022	180.54
IBM BG/P	PPC440	0.7	4	2.8	2	15	11.2	56	5.5	0.98	0.98	1.78	27	2600.6	9.97	1461	260.89
Tensilica	Custom	0.65	4	2.6	32	22	172.8	864	51.2	0.62	0.62	0.12	2.5	75	9.98	625	7.51
Movidius	Fragrak	0.8	6	4.8	128	2.8	204.8	1024	76.8	0.13	0.4	0.016	0.0455	3.25	9.98	200	0.33

http://www.hpcuserforum.com/presentations/Germany/EnergyandComputing _______Stgt.pdf

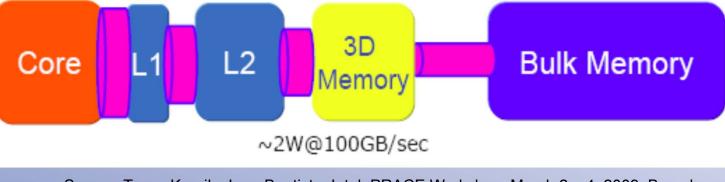
http://www.lbl.gov/cs/html/greenflash.html

http://www.tensilica.com/uploads/pdf/ieee_computer_nov09.pdf









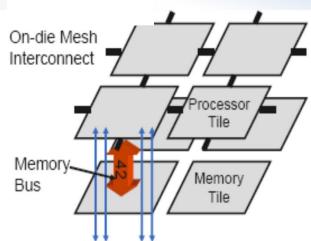
Source: Tanay Karnik, Jerry Bautista, Intel, PRACE Workshop, March 2 - 4, 2009, Barcelona





ADVANCED COMPUTING RESEARCH LABORATORY

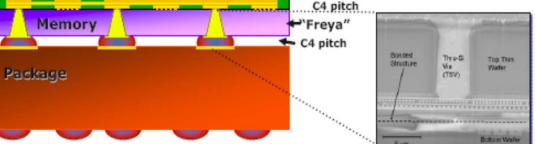
3D Memory Architecture



Signals and power from package, through memory, to the processor tile

TSV Pitch	190µm
SRAM die size	275mm ²
SRAM size	256KB per tile, 20MB total
SRAM Power	7W SRAM + 2.2W IO
Bandwidth	12GB/sec/tile, ~1TB/sec total

Work in Progress: Stacked Memory Prototype 256 KB SRAM per core 4X C4 bump density 3200 thru-silicon vias -tile processor with Cu bumps **Polaris*





Source: Tanay Karnik, Jerry Bautista, Intel, PRACE Workshop, March 2 – 4, 2009, Barcelona





DVANCED COMPUTING RESEARCH LABORATORY

HMC_{Gen1}: Technology Comparison

Generation 1 (4 + 1 memory configuration)

Technology	VDD	IDD	BW GB/s	Power (W)	mW/GB/s	pj/bit	real pJ/bit
SDRAM PC133 1GB Module	3.3	1.50	1.06	4.96	4664.97	583.12	762
DDR-333 1GB Module	2.5	2.19	2.66	5.48	2057.06	257.13	245
DDRII-667 2GB Module	1.8	2.88	5.34	5.18	971.51	121.44	139
DDR3-1333 2GB Module	1.5	3.68	10.66	5.52	517.63	64.70	52
DDR4-2667 4GB Module	1.2	5.50	21.34	6.60	309.34	38.67	39
HMC, 4 DRAM w/ Logic	1.2	9.23	128.00	11.08	86.53	10.82	13.7

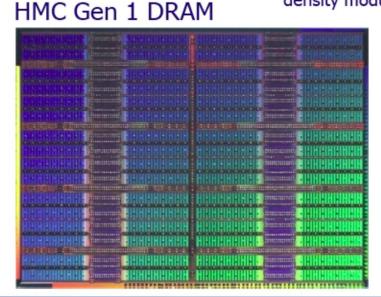
Simple calculation from IDD7 (SDRAM IDD4)

Real system, some with lower density modules

- 1Gb 50nm DRAM Array
- 90nm prototype logic
- 512MB total DRAM cube
- 128GB/s Bandwidth
- 27mm x 27mm prototype
- Functional demonstrations!

1Gb Based DRAM Stack

Reduced host CPU energy



Source: J. Thomas Pawlowski, Micron, HotChips23, August 17 - 19, 2011





ADVANCED COMPUTING RESEARCH LABORATORY

CHEP 2012 May 20 – 25, 2012 Lennart Johnsson

What type of Architecture?

Energy Efficiency – Embedded Processors

"A carefully designed ASIC can achieve an efficiency of 5 pJ/op in a 90-nm CMOS technology. In contrast, very efficient embedded processors and DSPs require about 250 pJ/op (50X more energy than an ASIC), and a popular laptop processor requires 20 nJ/op (4,000X more energy than an ASIC)"

Efficient Embedded Computing W. J. Dally, J. Balfour, D. Black-Shaffer, J. Chen, R. C. Harting, V.Parikh, J. Park, and D. Sheffield. 2008. Vol41, no.7, pp 27 – 32. IEEE Computer. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4563875&tag=1





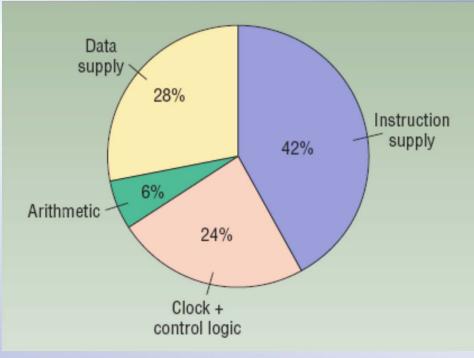


Advanced Computing Research Laboratory

Energy Efficiency – Embedded Processors

"An embedded processor spends most of its energy on instruction and data supply. The processor consumes 70 percent of the energy supplying data (28

percent) and instructions (42 percent). Performing arithmetic consumes only 6 percent. Of this, the processor spends only 59 percent on useful arithmetic— The operations the computation actually requires—with the balance spent on overhead, such as updating loop indices and calculating memory addresses."





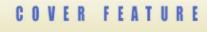
Efficient Embedded Computing W. J. Dally, J. Balfour, D. Black-Shaffer, J. Chen, R. C. Harting, V.Parikh, J. Park, and D. Sheffield. 2008. Vol41, no.7, pp 27 – 32. IEEE Computer. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4563875&tag=1





ADVANCED COMPUTING RESEARCH LABORATORY

Power Management



The Case for Energy-Proportional Computing

Luiz André Barroso and Urs Hölzle

Google



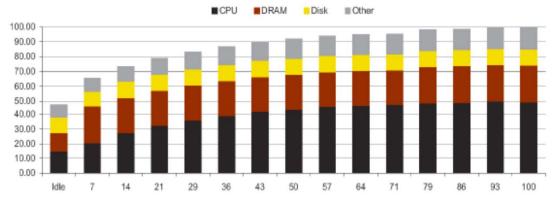
"The Case for Energy-Proportional Computing", Luiz André Barroso, Urs Hölzle, *IEEE Computer*, vol. 40 (2007).





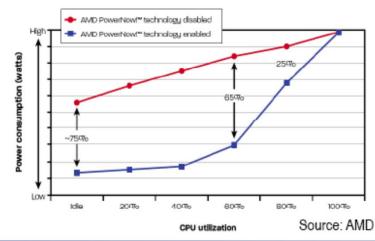
Power consumption - Load

Subsystem power usage varies from idle to full usage:



Font: Luiz Andre Barroso, Urs Hoelzle, "The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines", 2009.

Processors – Power stepping CPUs



CPUs can step down into reduced performance modes by adjusting frequency and voltage in synchronization with load.



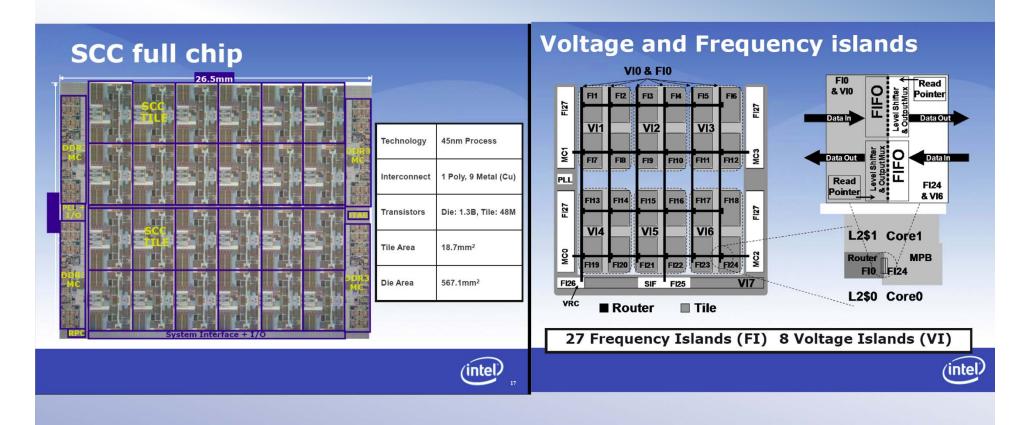
http://s3.amazonaws.com/ppt-download/greencomputing-2010-100202163517-phpapp02.pdf?Signature=G4jDh BFrZAkUbcbWo%2BnmKJafD78%3D&Expires=1281855097&AWSAccessKeyId=AKIAJLJT267DEGKZDHEQ





ADVANCED COMPUTING RESEARCH LABORATORY

Intel Single-chip Cloud Computer (2010)





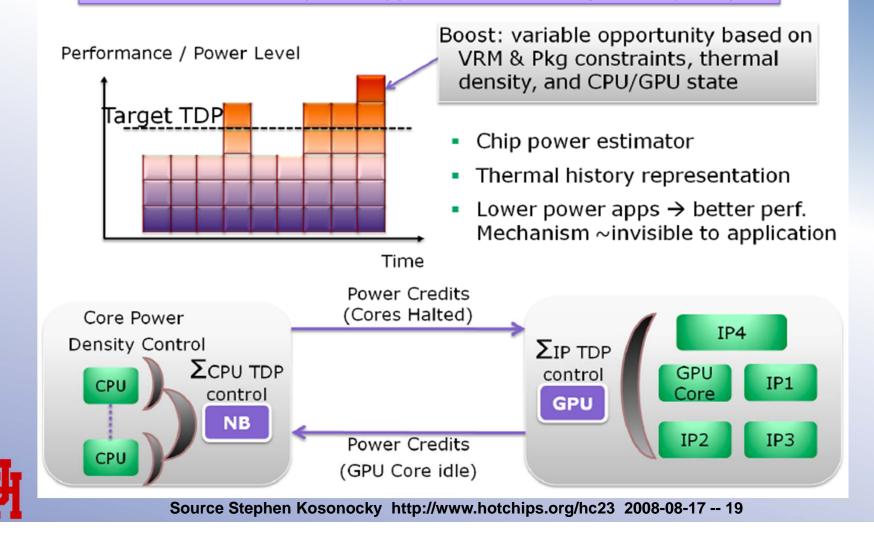
Source: Jim Held, Intel SCC Symposium February 12, 2010 http://techresearch.intel.com/newsdetail.aspx?ld=17#SCC





AMD Llano P-States

Performance: low-power applications run at higher frequency







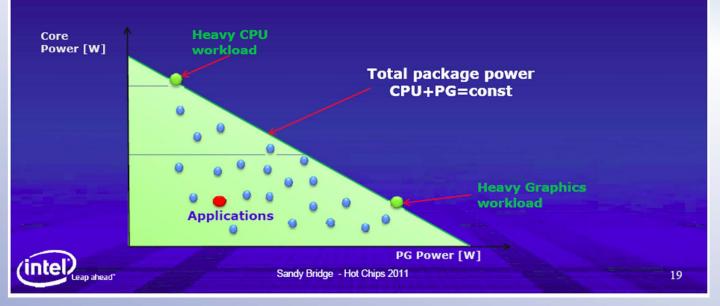
ADVANCED COMPUTING RESEARCH LABORATORY

Intel Sandy Bridge CPU

Intel® Turbo Boost Technology 2.0 - Package

Power specification is defined for the entire package

- Monolithic die power budget shared by CPU and PG
- Sum of component power at or below specifications



Source: Efi Rotem, Alon Naveh, Doron Rajwan, Avinash Ananthakrishnan, Eli Weissmann http://www.hotchips.org/hc23 2008-08-17 -- 19





ADVANCED COMPUTING RESEARCH LABORATORY

CHEP 2012 May 20 – 25, 2012 Lennart Johnsson



The Next (Final) Frontier?

Application Performance

Applications typically achieves ~1 - 10% of peak floating-point performance!!







Advanced Computing Research Laboratory

CHEP 2012 May 20 – 25, 2012 Lennart Johnsson



Thank You!

